Listing of Claims:

1. (Currently Amended) A method for allocating non-pageable pinned kernel memory for use by time sensitive data communications processing, the method comprising the steps of:

establishing a base non-pageable pinned kernel memory block;

accepting a request for a non-pageable pinned kernel memory buffer, wherein the request comprises a specification of a buffer size for the non-pageable pinned kernel memory buffer;

determining if the base non-pageable pinned kernel memory block contains sufficient non-pageable pinned kernel memory for the non-pageable pinned kernel memory buffer; and

allocating, in response to a determination that there is insufficient non-pageable pinned kernel memory within the base non-pageable pinned kernel memory block, an additional non-pageable pinned kernel memory block that is at least as large as the buffer size and wherein the additional non-pageable pinned kernel memory block is not required to form a contiguous non-pageable pinned kernel memory section with the base non-pageable pinned kernel memory block, and wherein the additional non-pageable pinned kernel memory block has a predetermined size.

2. (Previously Presented) The method according to claim 1, wherein the base non-pageable pinned kernel memory block, the additional non-pageable pinned kernel memory block and the non-pageable pinned kernel memory buffer are allocated by a non-pageable kernel memory allocation module.

Page 2 of 14

3. (Previously Presented) The method according to claim 1, wherein the additional non-pageable pinned kernel memory block is accessed through a linked list structure.

4. (Previously Presented) The method according to claim 1, wherein the additional non-pageable pinned kernel memory block has a predetermined size.

5. (Previously Presented) The method according to claim 1, wherein the request for the non-pageable pinned kernel memory buffer is received from a mass storage data server application processing module.

6. (Currently Amended) The method according to claim 1, further comprising the steps of:

     accepting a subsequent request for a second non-pageable pinned kernel memory buffer, wherein the subsequent request comprises a second specification of a second buffer size for the second non-pageable pinned kernel memory buffer;

     determining if the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block contain sufficient non-pageable pinned kernel memory for the second non-pageable pinned kernel memory buffer; and

     allocating, in response to a determination that there is insufficient non-pageable pinned ~~non-pageable~~ kernel memory within the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block, a second additional non-pageable pinned kernel memory block that is at least as large as the second buffer size and wherein the second additional non-pageable pinned kernel memory block is not required to form a contiguous non-pageable pinned kernel memory section with either the base non-pageable pinned kernel memory block or the additional non-pageable pinned kernel memory block.

7. (Previously Presented) A method of releasing memory allocated to a memory pool, wherein pinned memory buffers are temporarily allocated and not relocated within the memory pool and wherein additional memory blocks are sequentially added to the memory pool, the method comprising the steps of:

periodically examining a set of last two memory blocks to determine if at least one memory buffer is allocated therein; and

releasing a last added additional memory block if the step of periodically examining determines that there are no memory buffers allocated within the set of last two memory blocks.

8. (Currently Amended) A system for allocating non-pageable pinned kernel memory for use by time sensitive data communications processing, the system comprising:

a non-pageable pinned kernel memory; and

a non-pageable kernel memory allocation module, communicatively coupled to the non-pageable pinned kernel memory, wherein the non-pageable kernel memory allocation module:

establishes a base non-pageable pinned kernel memory block within the non-pageable pinned kernel memory;

accepts a request for a non-pageable pinned kernel memory buffer, wherein the request comprises a specification of a buffer size for the non-pageable pinned kernel memory buffer;

determines if the base non-pageable pinned kernel memory block contains sufficient non-pageable pinned kernel memory for the non-pageable pinned kernel memory buffer; and

allocates, in response to a determination that there is insufficient non-pageable pinned kernel memory within the base non-pageable pinned kernel memory block, an additional non-pageable pinned kernel memory block within the non-pageable pinned kernel memory that is at least as large as the buffer size, and wherein the additional non-pageable pinned kernel memory block has a predetermined size.

Page 4 of 14

9. (Previously Presented) The system according to claim 8, wherein the non-pageable kernel memory allocation module:

determines if the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block contain sufficient non-pageable pinned kernel memory for a second non-pageable pinned kernel memory buffer; and

allocates, in response to a determination that there is insufficient non-pageable pinned kernel memory within the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block, a second additional non-pageable pinned kernel memory block that is at least as large as the second buffer size and wherein the second additional non-pageable pinned kernel memory block is not required to form a contiguous non-pageable pinned kernel memory section with either the base pinned non-pageable kernel memory block or the additional non-pageable pinned kernel memory block.

10. (Previously Presented) The system according to claim 8, wherein the additional non-pageable pinned kernel memory block does not form a contiguous non-pageable pinned kernel memory block with the base non-pageable pinned kernel memory block.

11. (Previously Presented The system according to claim 8, wherein the additional non-pageable pinned kernel memory block is accessed through a linked list structure.

12. (Canceled)

13. (Previously Presented) The system according to claim 8, wherein the request for the non-pageable pinned kernel memory buffer is received from a mass storage data server application processing module.

14. (Previously Presented) A system for allocating memory for use by time sensitive data communications processing, the system comprising:

a kernel memory; and

a kernel memory allocation module, communicatively coupled to the kernel memory, wherein the kernel memory allocation module:

establishes a base memory block within the kernel memory;

accepts a request for a memory buffer, wherein the request comprises a specification of a buffer size for the memory buffer;

determines if the base memory block contains sufficient memory for the memory buffer; and

allocating, in response to a determination that there is insufficient memory within the base memory block, an additional memory block within the kernel memory that is at least as large as the buffer size; and

the system further comprising:

pinned memory buffers, wherein the pinned memory buffers are temporarily allocated and not relocated within a memory pool and wherein additional memory blocks are sequentially added to the memory pool, and wherein the kernel memory allocation module further:

periodically examines a set of last two memory blocks to determine if at least one memory buffer is allocated therein; and

releases a last added additional memory block if the step of periodically examining determines that there are no memory buffers allocated within the set of last two memory blocks.

15. (Previously Presented) The system according to claim 8, wherein the non-pageable kernel memory allocation module further:

accepts a subsequent request for a second non-pageable pinned kernel memory buffer, wherein the subsequent request comprises a second specification of a second buffer size for the second non-pageable pinned kernel memory buffer;

determines if the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block contain sufficient non-pageable pinned kernel memory for the non-pageable second pinned kernel memory buffer; and

allocates, in response to a determination that there is insufficient non-pageable pinned kernel memory within the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block, a second additional non-pageable pinned kernel memory block that is at least as large as the second buffer size.

16. (Currently Amended) A tangibly embodied computer readable medium including computer instructions for allocating non-pageable pinned kernel memory for use by time sensitive data communications processing, the computer instructions comprising instructions for:

establishing a base non-pageable pinned kernel memory block;

accepting a request for a non-pageable pinned kernel memory buffer, wherein the request comprises a specification of a buffer size for the non-pageable pinned kernel memory buffer;

determining if the base non-pageable pinned kernel memory block contains sufficient non-pageable pinned kernel memory for the non-pageable pinned kernel memory buffer; and

allocating, in response to a determination that there is insufficient non-pageable pinned kernel memory within the base non-pageable pinned kernel memory block, an additional non-pageable pinned kernel memory block that is at least as large as the buffer size and wherein the additional non-pageable pinned kernel memory block is not required to form a contiguous non-pageable pinned kernel memory section with the base non-pageable pinned kernel memory block, wherein the additional non-pageable pinned kernel memory block has a predetermined size.

17. (Previously Presented) The computer readable medium according to claim 16, wherein the base non-pageable pinned kernel memory block, additional non-pageable pinned kernel memory block and the non-pageable pinned kernel memory buffer are allocated by a non-pageable kernel memory allocation module.

18. (Previously Presented) The computer readable medium according to claim 16, wherein the additional non-pageable pinned kernel memory block is accessed through a linked list structure.

19. (Canceled)

Page 8 of 14

20. (Previously Presented) The computer readable medium according to claim 16, wherein the request for the non-pageable pinned kernel memory buffer is received from a mass storage data server application processing module.

21. (Previously Presented) The computer readable medium according to claim 16, further including computer instructions for:

accepting a subsequent request for a second non-pageable pinned kernel memory buffer, wherein the subsequent request comprises a second specification of a second buffer size for the second non-pageable pinned kernel memory buffer;

determining if the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block contain sufficient non-pageable pinned kernel memory for the second non-pageable pinned kernel memory buffer; and

allocating, in response to a determination that there is insufficient non-pageable pinned kernel memory within the base non-pageable pinned kernel memory block and the additional non-pageable pinned kernel memory block, a second non-pageable additional pinned kernel memory block that is at least as large as the second buffer size and wherein the second additional non-pageable pinned kernel memory block is not required to form a contiguous non-pageable pinned kernel memory section with either the base non-pageable pinned kernel memory block or the additional non-pageable pinned kernel memory block.

Page 9 of 14

22. (Currently Amended) A tangibly embodied computer readable medium including computer instructions for releasing memory allocated to a memory pool, wherein pinned memory buffers are temporarily allocated and not relocated within the memory pool and wherein additional memory blocks are sequentially added to the memory pool, the computer instructions comprising instructions for:

periodically examining a set of last two memory blocks to determine if at least one memory buffer is allocated therein; and

releasing a last added additional memory block if the step of periodically examining determines that there are no memory buffers allocated within the set of last two memory blocks.